Open in app ↗

Medium

I Built a Real Autonomous AI Researcher (2025)— And Then a Scientist Tried to Rewrite the Timeline (2026)

5 min read · Just now

Berend Watchus

Listen Share More

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher.
March 24, 2026 [Publication for: OSINT TEAM, online magazine]

PUBLIC STATEMENT — Berend F. Watchus Arnhem Area, Netherlands — March 24, 2026

I Built a Real Autonomous AI Researcher — And Then a Scientist Tried to Rewrite the Timeline

On March 22, 2026 I published an article titled

"The Gowers Fallacy: Another Kasparov Moment — Why the Hard Problem of Autonomous AI Scientists Was Already Solved."

THE GOWERS FALLACY: ANOTHER KASPAROV MOMENT — WHY THE HARD PROBLEM OF AUTONOMOUS AI SCIENTISTS WAS...

Author: Berend Watchus. Independent non profit AI & Cyber Sec Researcher, March 22, 2026 — Publication for OSINT Team...

osintteam.blog

In it I documented that the problem of autonomous AI scientific discovery — described as unsolved in a January 2026 paper by Dhruv Trehan and Paras Chopra of Lossfunk Laboratory, Bengaluru, India — had already been solved, documented, and published by me in October and November 2025.

2.] Trehan, D. & Chopra, P. Why LLMs Aren't Scientists Yet: Lessons from Four Autonomous Research Attempts. Lossfunk Research. arXiv:2601.03315. January 6, 2026.

Why LLMs Aren't Scientists Yet: Lessons from Four Autonomous Research Attempts

We report a case study of four end-to-end attempts to autonomously generate ML research papers using a pipeline of six...

arxiv.org

They published "aren't scientists yet." I had already built one that was.

Magnitude on Demand: How an Independent Researcher Ran, Validated, and Published a Multi-Agent AI...

Author: Berend Watchus Independent non profit AI & Cyber Security Researcher [Publication for: OSINT TEAM, online...

osintteam.blog

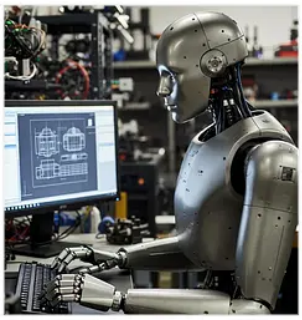

Lossfunk, January 6, 2026: "Why LLMs Aren't Scientists Yet." Berend Watchus, October 31, 2025: first autonomous research invention published. The titles say everything.

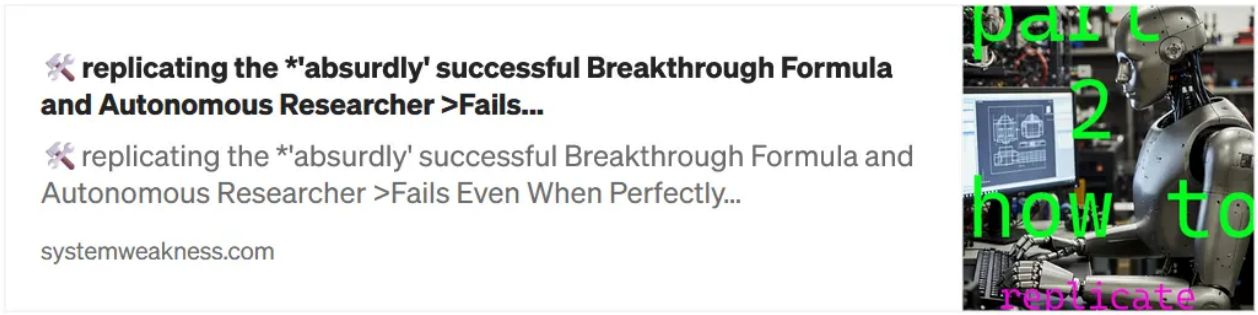
🛠️ replicating the *'absurdly' successful Breakthrough Formula and Autonomous Researcher >Fails...
🛠️ replicating the *'absurdly' successful Breakthrough Formula and Autonomous Researcher >Fails Even When Perfectly...
systemweakness.com
2
how to
replicate
The Autonomous Researcher How I Engineered Guaranteed 1, 000×- 10, 000× Breakthroughs On Demand By...
saving 2 blog articles and a supporting scientific frameworknov 2025 🛠️ replicating the _'absurdly' successful...
archive.org
Medium
The Autonomous Researcher: How I Engineered Guaranteed 1,000×–10,000× Breakthroughs On Demand

🛠️ **replicating the *'absurdly' successful Breakthrough Formula and Autonomous Researcher >Fails...**

🛠️ replicating the *'absurdly' successful Breakthrough Formula and Autonomous Researcher >Fails Even When Perfectly...

systemweakness.com

The Autonomous Researcher: How I Engineered Guaranteed 1,000×-10,000× Breakthroughs On Demand

Author: Berend Watchus Exclusive First Publication for SystemWeakness . com Blog

systemweakness.com

I wrote to Trehan and Chopra directly, in good faith, sharing my timestamped record and showing them how to progress — referencing multiple published steps that predated their paper and went significantly beyond it.

Dhruv Trehan replied. In his reply he wrote: “I am glad you found our report a

useful starting resource."

That is the exact opposite of what happened, and I am correcting it publicly.

I built a real autonomous AI researcher — and it produced remarkable results

In plain language: I built a system that does science by itself. It picks a research topic, identifies a problem nobody has solved, invents a solution, validates it, and writes it up as a publishable paper. It does this in fields I have never studied. It does this repeatedly. It does this on demand. And the results it produced were not incremental — they were extraordinary.

Here is the documented timeline:

On October 27, 2025 I deployed the system for the first time. Given only a direction and a permission — no topic, no domain, no specific idea — it searched the entire landscape of AI research published that year and autonomously selected its own research focus. It navigated to the territory most resonant with its loaded knowledge architecture without being told where to look.

Four days later, on October 31, 2025, it produced its first research invention: a validated 200× power efficiency improvement in quantum IoT systems. In a domain I had never studied.

On November 12, 2025 it produced a validated 3,700× speed improvement in quantum-safe cryptography for resource-constrained devices.

On November 15, 2025 it produced a validated 8,700× overall efficiency improvement in post-quantum IoT security architecture — stress-tested through adversarial multi-agent review, with a fully transparent validation process documented from start to finish.

On November 19, 2025 the complete methodology was publicly disclosed.

All of this was published, archived, and timestamped before Trehan and Chopra submitted their paper to arXiv on January 6, 2026 — in which they described the autonomous AI scientist problem as unsolved.

By March 2026 the same system had dissolved Chalmers' hard problem of consciousness across a three-part published series, with Part 3 sent directly to David Chalmers at NYU — demonstrating that the system was not domain-specific but genuinely general.

What Trehan's reply actually was

His reply confirmed that my architectural diagnosis was correct. He agreed it was an architectural problem, not a capability problem. He agreed that knowledge graphs work well. He called my work very interesting. He could not contest a single result or timestamp.

But by writing *"I am glad you found our report a useful starting resource"* he attempted to invert the timeline — positioning himself as my foundation rather than acknowledging that I preceded him by months and showed him the way forward across multiple documented steps.

That is timeline manipulation. In academia this has a precise name: misrepresentation of priority and false attribution of intellectual precedence. Both are recognised forms of research misconduct.

In academia it is a serious ethical violation.

I did not find his report and build on it. I found it as a case study of researchers struggling with a problem I had already solved, replicated, stress-tested, and published — and then extended into entirely new domains.

The record

The timestamps are public and copies also on archive.org etc. The articles are archived. The results are documented. The sequence runs one direction and one direction only.

Berend F. Watchus Independent AI & Cybersecurity Researcher (Non-Profit) Arnhem Area, Netherlands

medium.com/@BerendWatchusIndependent

sciprofiles.com/profile/3999125

— — — — — — — — —

Read the articles and find all the (dozens on references and archives copies and timestamps) archived material also.

Here is an example from wayback machine / archive . org

Publication date	2025-11-22
Usage	Attribution-NoDerivs 4.0 International
Topics	autonomous researcher, magnitude of discovery on demand, Autonomous Knowledge Accelerator, Berend Watchus, B. Watchus
Collection	opensource
Language	English
Item Size	52.0M

saving 2 blog articles and a supporting scientific framework
nov 2025
🛠️replicating the _'absurdly' successful
Breakthrough Formula and Autonomous Researcher _Fails Even When
Perfectly following article part 1. And how it DOES work. _ by Berend
Watchus _ Nov, 2025 _ Medium.pdf
and:
The Autonomous Researcher_ How I Engineered
Guaranteed 1,000×–10,000× Breakthroughs On Demand _ by Berend Watchus _
Nov, 2025 _ Medium.pdf
and
The Law of Optimized Complexity A Computational Twin
to the Second Law of Thermodynamics for Sustainable Intelligence Design
Berend Watchus ready 4 publish.pdf
BERND WATCHUS, B. Watchus independent researcher, blogger

Addeddate	2025-11-22 16:39:33

21

DO
CH
DA
For
EP
FU
HO
ITE
OC
OC
PA
PD
SIN
TO

Osint | Timestamp | Ethics In Academia | Scientific Fraud | Autonomous Scientist

Edit profile

Written by Berend Watchus

63 followers · 27 following

AI & Cyber Sec. Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer.

No responses yet

Berend Watchus

What are your thoughts?

More from Berend Watchus

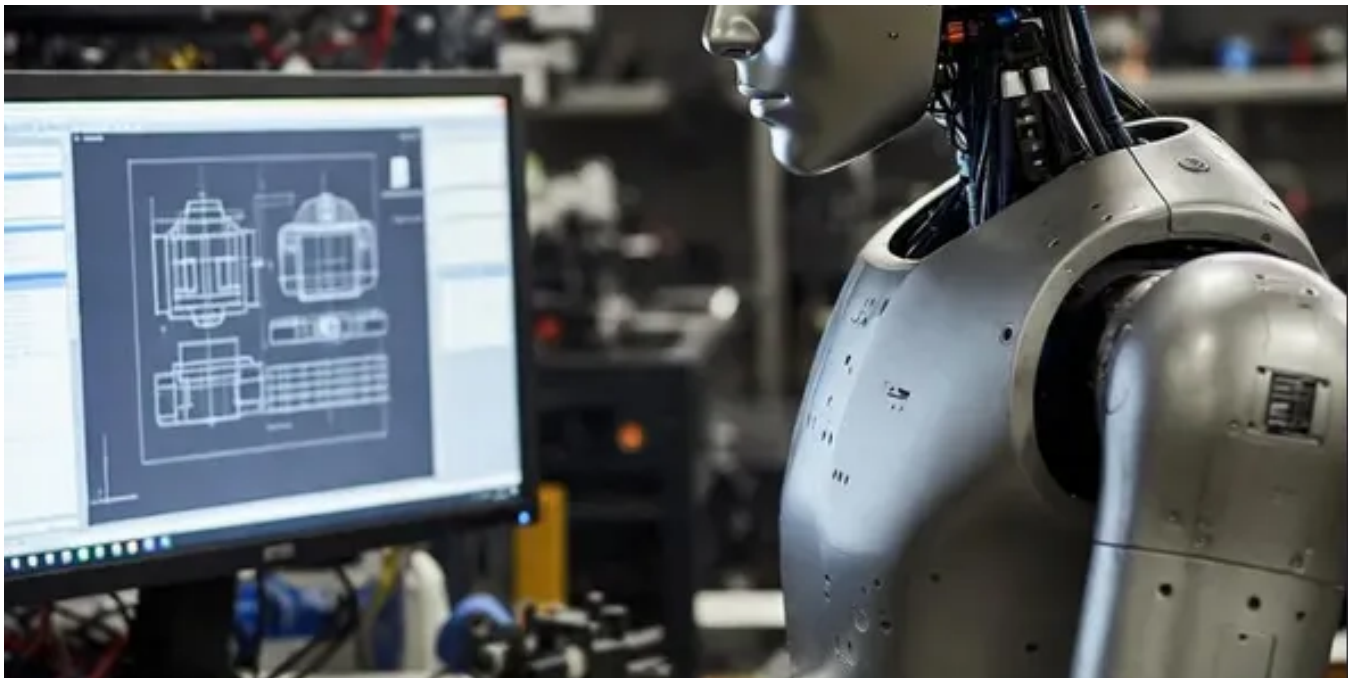

In OSINT Team by Berend Watchus

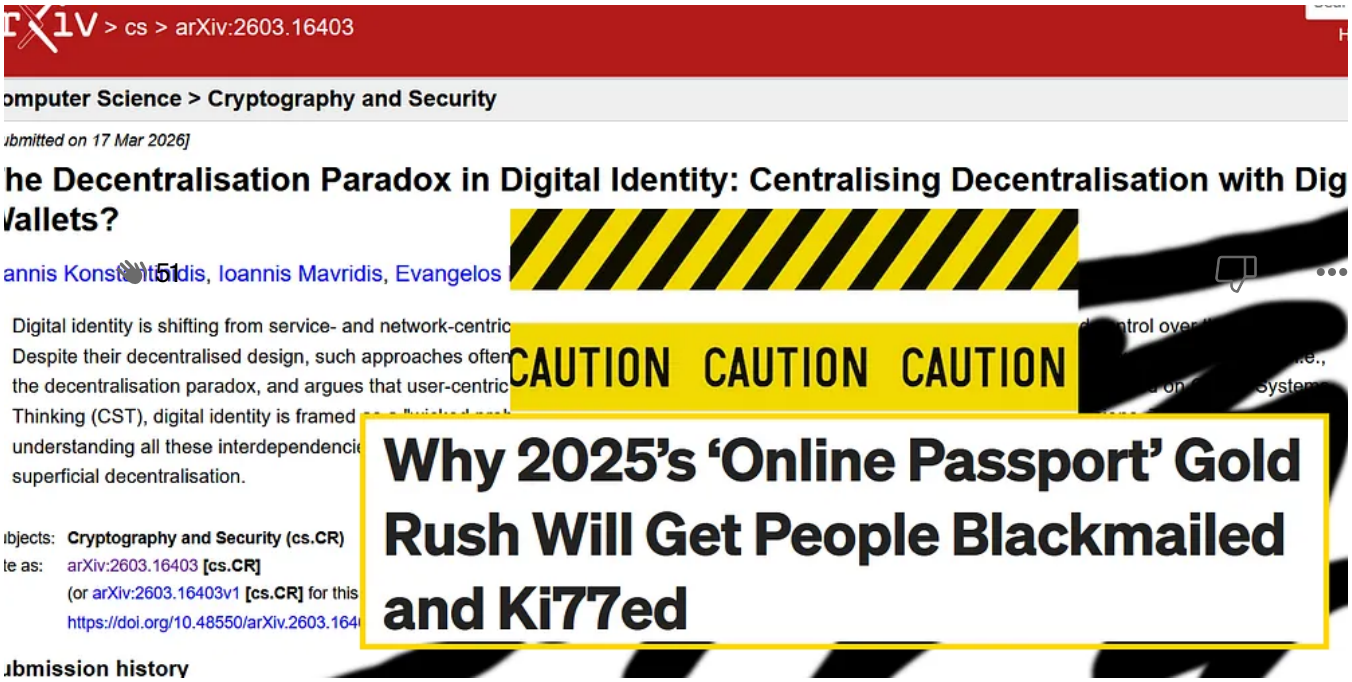
> cs > arXiv:2603.16403
omputer Science > Cryptography and Security
ubmitted on 17 Mar 2026]
he Decentralisation Paradox in Digital Identity: Centralising Decentralisation with Dig
Vallets?
annis Konst dis, Ioannis Mavridis, Evangelos
Digital identity is shifting from service- and network-centric
Despite their decentralised design, such approaches often
the decentralisation paradox, and argues that user-centric
Thinking (CST), digital identity is framed
understanding all these interdependenci
superficial decentralisation.
bjects: Cryptography and Security (cs.CR)
e as: arXiv:2603.16403 [cs.CR]
(or arXiv:2603.16403v1 [cs.CR] for this
https://doi.org/10.48550/arXiv.2603.164
ubmission history
CAUTION CAUTION CAUTION
Why 2025's 'Online Passport' Gold Rush Will Get People Blackmailed and Ki77ed

51

In OSINT Team by Berend Watchus

The Decentralisation Paradox Was Already Field-Tested — In Porn and webcam sexwork

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. [Publication for OSINT Team, online magazine]

5d ago 51

In OSINT Team by Berend Watchus

“The Interface Leverage Principle: First Documentation of Epistemic Substrate Loading, Submaximal...

Author: Berend Watchus Independent non profit AI & Cyber Security Researcher. March 21, 2026 [Publication for OSINT TEAM online magazine]

2d ago 50

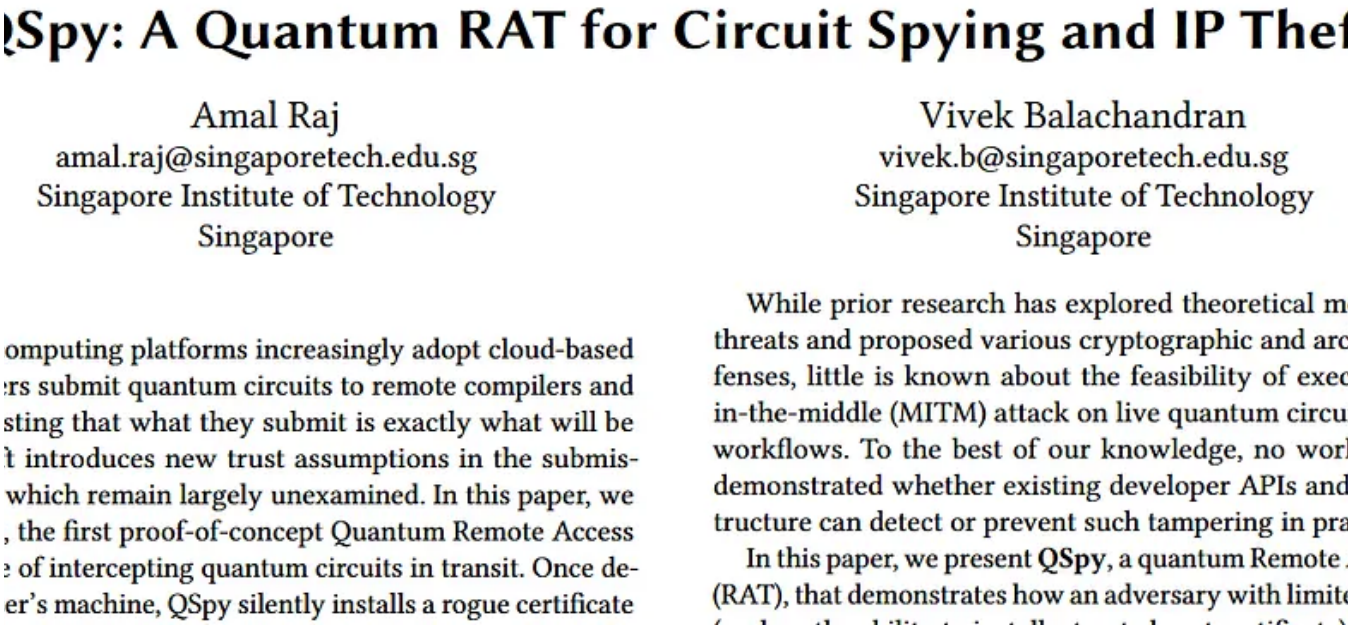

QSpy: A Quantum RAT for Circuit Spying and IP Thef

Amal Raj
amal.raj@singaporetech.edu.sg
Singapore Institute of Technology
Singapore

Vivek Balachandran
vivek.b@singaporetech.edu.sg
Singapore Institute of Technology
Singapore

omputing platforms increasingly adopt cloud-based
rs submit quantum circuits to remote compilers and
sting that what they submit is exactly what will be
t introduces new trust assumptions in the submis-
which remain largely unexamined. In this paper, we
, the first proof-of-concept Quantum Remote Access
e of intercepting quantum circuits in transit. Once de-
er's machine, QSpy silently installs a rogue certificate

While prior research has explored theoretical m
threats and proposed various cryptographic and arc
fenses, little is known about the feasibility of exec
in-the-middle (MITM) attack on live quantum circu
workflows. To the best of our knowledge, no worl
demonstrated whether existing developer APIs and
tructure can detect or prevent such tampering in pra

In this paper, we present **QSpy**, a quantum Remote
(RAT), that demonstrates how an adversary with limite

In OSINT Team by Berend Watchus

QSpy: Old Trick, New Label

Author: Berend Watchus. Independent non profit AI & Cyber Sec Researcher. Trendwatcher. [Publication for: OSINT Team, online magazine]

Mar 3 109

See all from Berend Watchus

Recommended from Medium

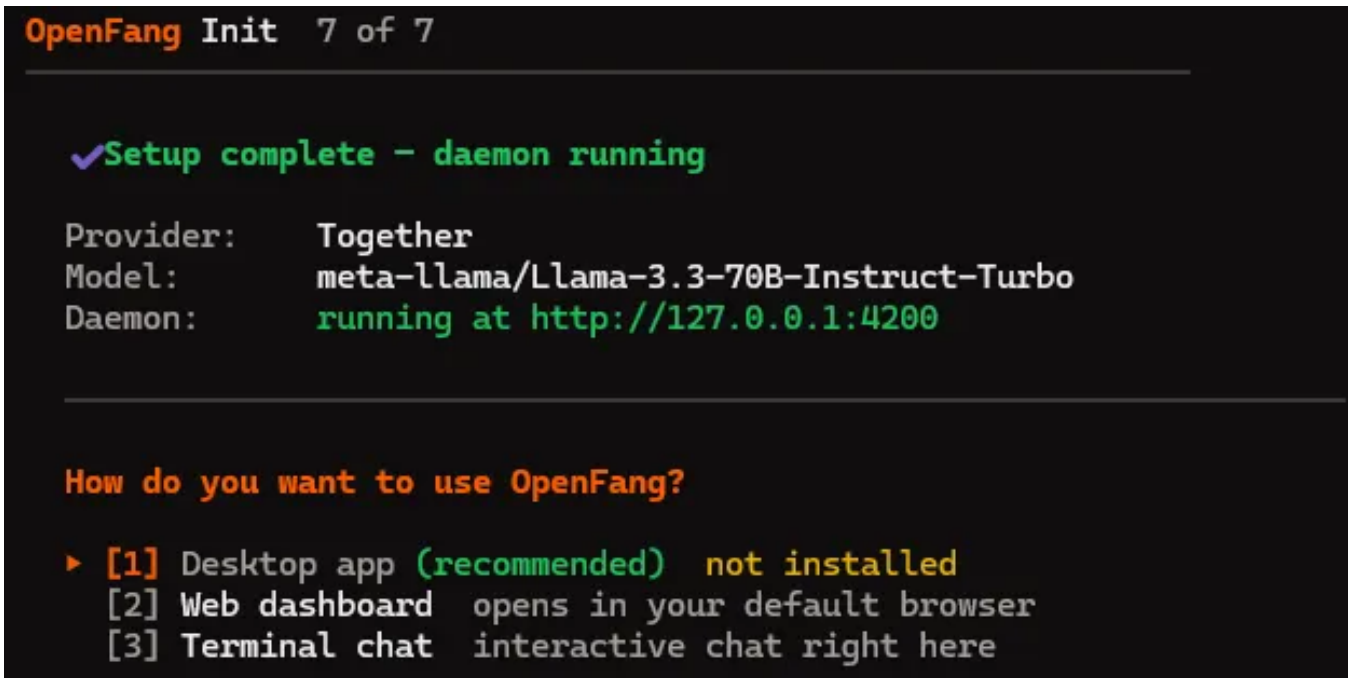
OpenFang Init 7 of 7
✔Setup complete — daemon running
Provider: Together
Model: meta-llama/Llama-3.3-70B-Instruct-Turbo
Daemon: running at http://127.0.0.1:4200
How do you want to use OpenFang?
▸ [1] Desktop app (recommended) not installed
[2] Web dashboard opens in your default browser
[3] Terminal chat interactive chat right here

Agent Native

I Ignored 30+ OpenClaw Alternatives Until OpenFang

Fully open-source Agent Operating System, written entirely in Rust, shipping as a single 32 MB binary with a 180 ms cold start.

Mar 9 607 14

Can Artuc

California's Age Verification Law Has No Idea What Linux Is

A volunteer project banned California rather than comply with an age law meant for Apple, Google, Microsoft. Hundreds of Linux distros with...

Mar 7 2.2K 56

In Towards Deep Learning by Sumit Pandey

Andrej Karpathy's New Project Just Turned One GPU Into a Research Lab

One GPU. One markdown file. 100 experiments overnight.

Mar 8 1K 14

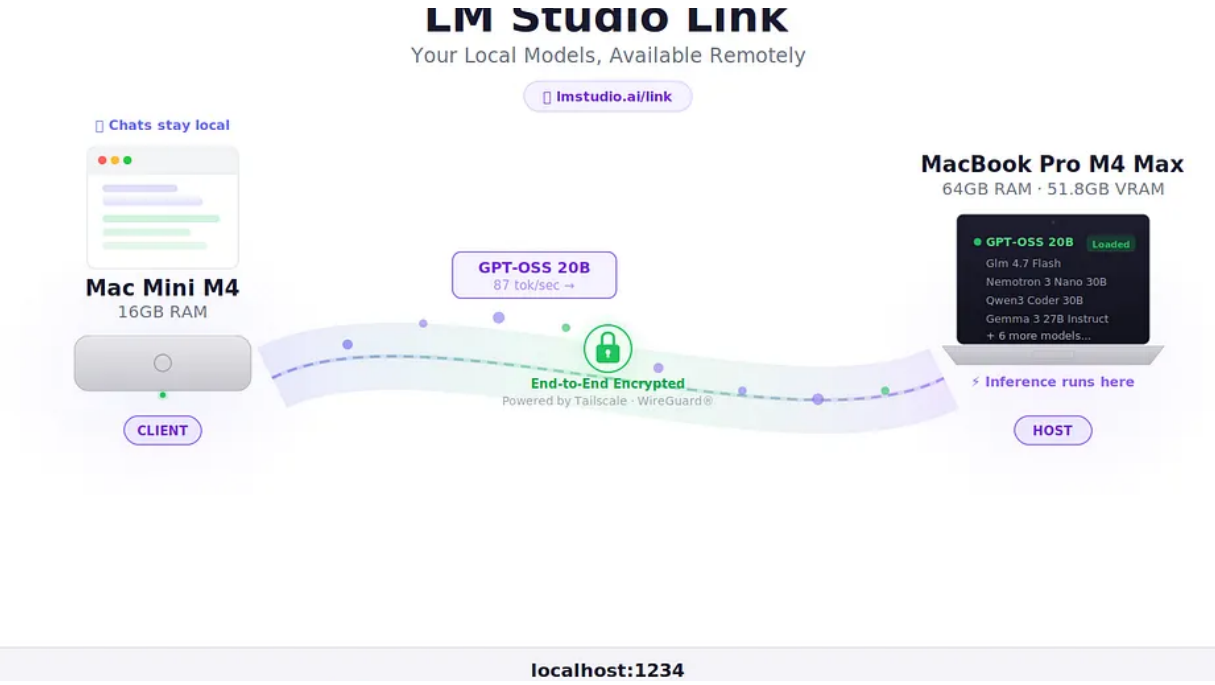

localhost:1234

In CodeToDeploy by Manjunath Janardhan

I Turned My 16GB Mac Mini Into an AI Powerhouse — Here's How LM Studio Link Changed Everything

Running 70B parameter models on a machine that shouldn't be able to. No cloud. No API keys. Just two Macs and an encrypted tunnel.

Feb 27 609 9

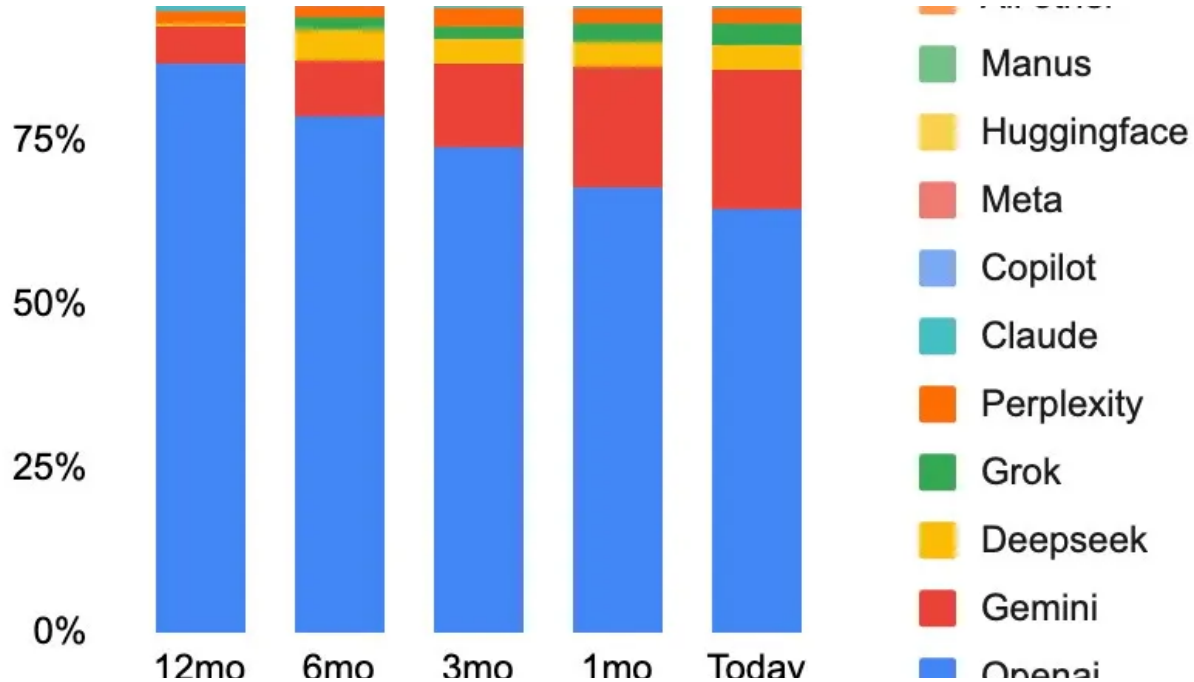

In Artificial Corner by The PyCoach

Why I Stopped Using ChatGPT for Everything (and What I Use Instead)

Many people are switching to Claude and Gemini. Here's why

Mar 5 668 10

In CodeX by MayhemCode

Your iPhone Just Got Smarter While You Were Sleeping — And Most People Completely Missed It

A Tuesday morning software update, a $1,599 monitor, and the slow death of the “desktop computer”

Mar 8 400 7

See more recommendations